Cyberstability Paper Series

**New Conditions and Constellations in Cyber**

# Routing Without Rumor

## Securing the Internet's Routing System

**Danny McPherson**
Executive Vice President & Chief Security Officer, Verisign

December 2021

Cyberstability Paper Series
**New Conditions and Constellations in Cyber**

# Routing Without Rumor: Securing the Internet's Routing System

**Danny McPherson** | Executive Vice President & Chief Security Officer, Verisign

December 2021

The Global Commission on the Stability of Cyberspace (GCSC) has spent a considerable amount of time and resources developing eight norms by which to influence state and non-state behaviors to support the stability of cyberspace.[1] One of these norms focuses on "the public core of the Internet," which at a high level constitutes "such critical elements of the infrastructure of the Internet as packet routing and forwarding, naming and numbering systems, the cryptographic mechanisms of security and identity, transmission media, software, and data centers." A more detailed definition of the Norm on the Non-interference with the Public Core[2] is available on the GCSC website.

The Norm declares that "State and non-state actors should neither conduct nor knowingly allow activity that intentionally and substantially damages the general availability or integrity of the public core of the Internet, and therefore the stability of cyberspace."

This paper, in the GCSC's "New Conditions and Constellations in Cyber" Cyberstability Paper Series,[3] is primarily concerned with the public core of the Internet's packet routing and forwarding elements, as well as with corresponding Internet numbering systems. We'll first provide some background information on the Internet architecture and Internet number resource allocation, and then discuss some vulnerabilities in the Internet routing system and what mechanisms are aiming to mitigate those vulnerabilities. We'll then provide some considerations all stakeholders need to consider as we aim to find a balance between vital new infrastructure components, such as Resource Public Key Infrastructure (RPKI) that aims to help secure the routing system, and the implications that come along with its adoption.

**Danny McPherson** is Executive Vice President and Chief Security Officer at Verisign. He has authored several books, numerous internet protocol standards, network and security research papers, and other publications.

The opinions expressed in this publication are those solely of the author(s) and do not reflect the views of the Global Commission on the Stability of Cyberspace (GCSC), its partners, or The Hague Centre for Strategic Studies (HCSS).

© 2021 The Hague Centre for Strategic Studies and the Global Commission on the Stability of Cyberspace. This work is licensed under a Creative Commons Attribution –Noncommercial – No Derivatives License.

## The Internet Protocol (IP) and Internet Number Resources

The Internet is made up of a loosely interconnected network of networks. These networks utilize the Internet Protocol (IP) suite, a collection of technical standards and rules, to relay packets within and between networks. IP provides the formatting of data exchanged as well as the addressing system, and a routing function is provided by systems referred to as routers that enables the inter-networking, allowing information to be exchanged between networks and creating a unified single global network—the Internet.

IP addresses are used to uniquely identify each device on the Internet. There are two types of IP addresses used on the Internet today: the 32-bit IP version 4 (IPv4) addresses, which allow for unique addresses of just over ~4 billion endpoints (232), which seemed sufficient when the Internet was first developed, and a newer version of IP, the 128-bit IP version 6 (IPv6) addresses, which provides ~340 undecillion (2128) of available addresses.

Just as with phone numbers, global uniqueness of IP addresses for devices connected to the Internet is crucial. To maintain uniqueness of IP addresses, global coordination and allocation is required. As illustrated in Figure 1, IP addresses are distributed in blocks (i.e., address ranges) from the Internet Assigned Numbers Authority (IANA) to the five Regional Internet Registries (RIRs), who assign them to National Internet Registries (NIRs), Local Internet Registries (LIRs), or directly to Internet Service Providers (ISPs) and end users.[4]

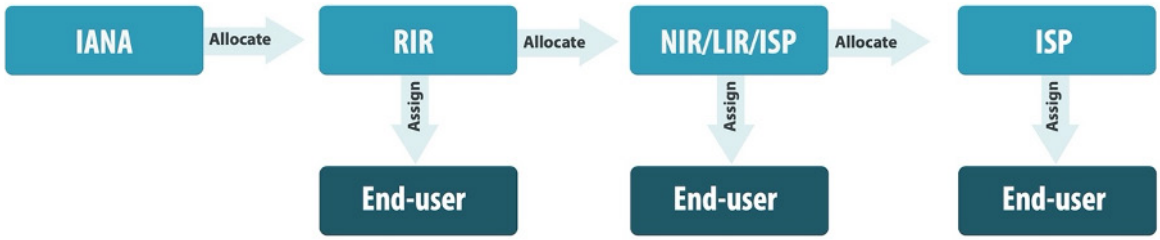


**Figure 1: Internet Number Resource Allocation: IP Addresses and AS Numbers (source: ARIN.net)**

In addition to IP addresses, each network that connects to the Internet needs to obtain a unique Autonomous System (AS) number, which is used by routing protocols to identify that network within the global routing system. These AS numbers are distributed in the same manner as IP addresses. AS numbers were originally specified as 16-bits, allowing for AS numbers from 0 through 65535. In the mid-2000s the Internet Engineering Task Force (IETF)[5] developed backwards-compatible 32-bit AS numbers (~4 billion) and transitioned to the larger AS numbers. Today, AS numbers are allocated from this larger number space, and it's a good thing, given that there are already ~72,500 unique ASes represented in the global routing system currently.[6]

The collection of network devices, border and internal routers that comprise each network connected to the Internet vary considerably. For example, a small enterprise may only have one low-end internal and Internet-connect router, whereas a large enterprise, regional ISP, or university may have hundreds or thousands, and a large ISP may have thousands or even tens of thousands of routers.

Similarly, where these networks connect to the Internet will vary. Small enterprises may only connect in one location to a regional or local ISP, whereas large enterprises may connect in tens or

hundreds of locations, and interconnect with other networks either directly or at one or more Internet Exchange Points (IXPs).[7] Large ISPs may interconnect with other networks in multiple locations and across many regions and countries, as well as via a multitude of IXPs. Regardless of where and how they interconnect, if they're connecting to and participating in the global routing system, they'll generally use a single AS number to uniquely identify their network. Each individual network is designed to support the business and policy objectives of that individual network's administrators. There is no centralized planning authority or coordination facility dictating how or where networks interconnect globally.

Correspondingly, the number of network administrators will vary considerably, where there may be only one or two at a small network, but potentially hundreds at a large ISP. In aggregate, there may be a million or more individuals involved with routing on the global Internet.

## Internet Routing and the Border Gateway Protocol (BGP)

Networks often interconnect at a multitude of locations. The primary job of the routing system is to learn all available paths through the network(s) to reach a particular destination, and when faced with a multitude of paths for a given route, to use what local administrators deem as the best route at any given instant. In the routing system, these destinations are codified as blocks of IP addresses, commonly referred to as prefixes (much like the telephone numbering system), and metadata is added to the prefix identifying the network(s) the information traversed within the routing system to reach the local router. This prefix and associated metadata constitute what's referred to as destination network layer reachability information, or a "route." Routes can be for either IPv4 or IPv6 destinations, and there are ~903,000 IPv4 prefixes in the current routing system, and ~142,000 IPv6 prefixes.[8]

Time and again, the Internet routing system has proven to be highly effective and robust in the face of localized and regional failures, finding alternative available routes to a destination if the current preferred route becomes unavailable. The global routing system has dealt with immense scaling challenges across multiple dimensions (e.g., the number of ASes, the number of discrete interconnections, the growth of routes, the number of available paths to reach a given destination, and the amount of instability or "churn" in the system).

Today, the routing system largely relies on a decentralized and implicit trust model of network self-assertions that effectively creates a transitive "web of trust." There is no central authority dictating which networks are authorized to assert reachability for an Internet destination.

The Border Gateway Protocol (BGP),[9] standardized by the IETF, is the de-facto inter-domain routing protocol on the Internet. Conceptually, border routers within each AS establish BGP peering sessions internally, as well as across each point of interconnection with border routers in other ASes, and the routers are referred to as BGP neighbors, as illustrated in Figure 2. In accordance with local routing policies, each router advertises destination reachability information to each of their BGP neighbors, effectively self-asserting that they provide reachability to the collections of IP addresses within the IP address block(s) represented by the route(s). It is these routing policies that therefore decide where and how Internet traffic flows, which not only factors into account performance characteristics, such as availability and latency, but also potentially the security of resulting data that will be transmitted, as well as the financial cost of exchanging data in certain locations. Understanding how routing works is therefore a major factor in understanding both Internet security and Internet economics.

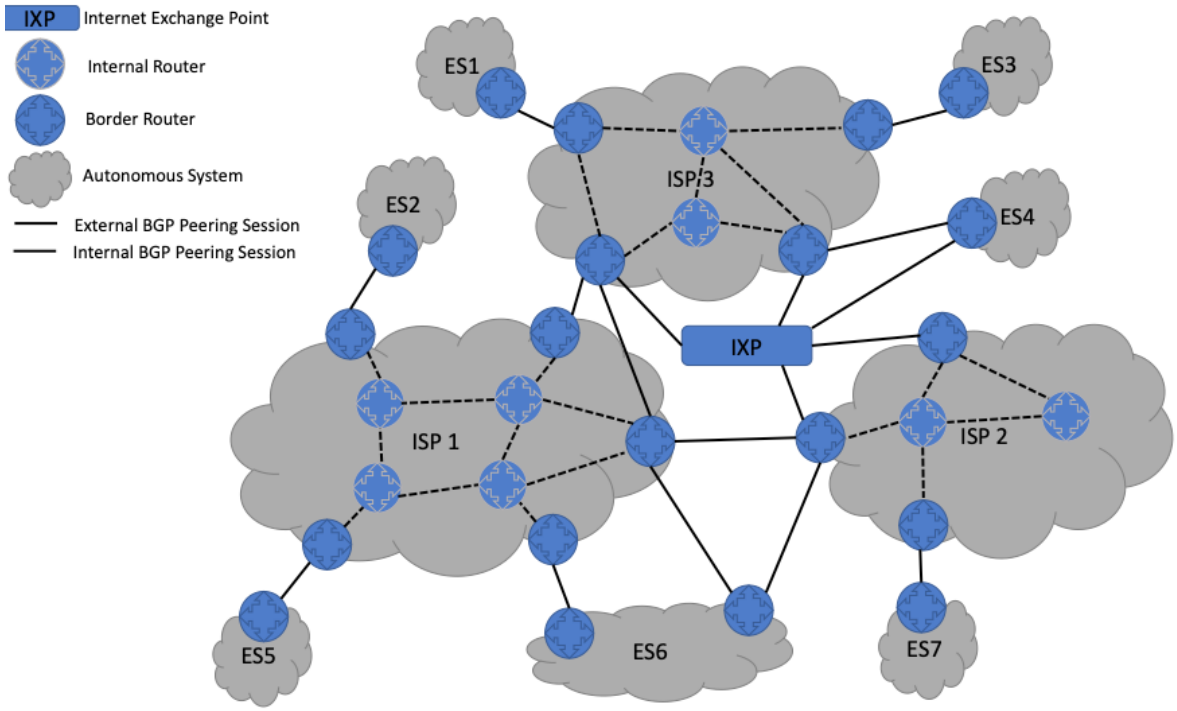


**Figure 2: Sample Inter-domain Interconnection Model**

Today, the routing system largely relies on a decentralized and implicit trust model of network self-assertions that effectively creates a transitive "web of trust." There is no central authority dictating which networks are authorized to assert reachability for an Internet destination. Each individual AS independently applies its own locally provisioned policies, choosing what action to take on each of the destinations for which it locally provides connectivity, as well as on all of the routes it received from other networks. For each route received from a peer, a router may choose to

1. only use the information locally for packet forwarding (e.g., in Figure 2 if ISP1 were to receive a route for a destination connected to End Site 4 (ES4) from ISP3, they might choose to only share it with end sites (customers) connected to them, and not with ISP2, or
2. use the information locally as the preferred route and propagate it (i.e., advertise destination reachability) to one or more of its peer networks, to include ISP2, which could result in ISP1 being in the datapath for the route if ISP2 has no other route, or
3. simply discard or suppress the route received from the peer and not share it with anyone.

When a preferred route to a destination learned via a given path (e.g., internally or via a BGP neighbor in another AS, as illustrated in Figure 2) becomes unavailable, and if an alternative path via another router exists, the alternative can immediately be used to reprogram the router's packet forwarding logic and the router can continue to transmit traffic toward the destination. This may result in a less desirable path being used, e.g., in action 3 above where traffic may flow from ISP2 through ISP1 to get to ISP3 and ultimately, ES4 if the IXP were to become unavailable.

Despite being designed over three decades ago in a vastly different Internet, BGP has scaled so well because (a) it operates in a distributed manner, (b) it has no central point of control and therefore of failure, and (c) each network acts autonomously with regard to whom it interconnects with

and what information it chooses to use and/or propagate. While an array of pricing, performance, and security characteristics are used to develop routing policies in each AS, ultimately BGP will use any available path to reach a destination, and often enough the choice of how to route between two ASes is dependent upon interpersonal factors between the individual network administrators themselves, and upon informal assessments of technical and even personal reliability—this behavior could be considered routing by rumor. In the idealized scenario where network operators only deal with noble actors, and none of the million(s) of network administrators are capable of mistakes, and there is zero probability that bad actors would gain access to one or more of those networks, then this distributed system would function well and could be fully trusted. But pragmatically in today's world, where routing incidents continue to cause operational and security issues, operators know the idealized scenario is not the case. Much like the Domain Name System (DNS) and other early Internet infrastructure protocols in which ease of use, open end-to-end connectivity, system resilience, and scalability were primary objectives, security was an afterthought.

Much like the Domain Name System (DNS) and other early Internet infrastructure protocols in which ease of use, open end-to-end connectivity, system resilience, and scalability were primary objectives, security was an afterthought.

## Routing Security Incidents

The two most prominent types of operational and security incidents that occur in the routing system today are "route hijacks"[10] and "route leaks."[11] Route hijacks involve the accidental or malicious rerouting of internet traffic and are sometimes referred to as mis-origination,[12] in which the originating AS contained within the BGP metadata associated with the route is usually not the legitimate origin. Route leaks typically involve the unintentional or malicious propagation of routing information beyond the intended scope of the originator, receiver, and/or one of the networks along the route's path, thereby resulting in potentially unintended or undesirable networks being inserted into the datapath used to reach a given destination. For example, imagine a scenario where ES6 in Figure 2 began announcing to ISP2 routes that it had learned from ISP1, and because routing policies commonly prefer customer-learned routes over peer-learned routes, traffic from ISP2 to ISP1 destinations begins flowing through ES6. As a result, there could be latency and packet loss issues, as well as potential man-in-the-middle (MitM) or denial of service attack conditions as a result.

Often, route leaks will preserve the originating AS in the route's metadata, but that's not always the case. Interestingly, if the origin is preserved during a route leak, then many of the origin validation controls that may be in place are implicitly circumvented. Both route hijacks (e.g., how Pakistan Telecom effectively globally exported their state censorship on YouTube services[13]) and route leaks[14] can result in partial or full rerouting of traffic for the impacted destinations. This can potentially result in changes to the packet forwarding path and have an array of security implications. These include enabling denial of service conditions, when traffic is selective or discarded wholesale either intentionally or because of insufficient resources to forward it on to the intended destination. It can also facilitate man in the middle (MitM) and Man-on-the-Side, and other "on-path" attacks, which can allow an attacker the opportunity with which to influence the confidentiality, availability, and even integrity of the data stream, depending on the attacker's sophistication and the type of encryption used. Not all the incidents are intentional. However, discerning intent is extremely difficult given that the complexity of routing policy configuration, deployment, and implementation vary considerably.

Some routing incidents may also simply be a misconfiguration that results in added latency, and/or potential network congestion, in reaching the destination without attack exposure. Leaks occur very frequently in the routing system[15] and it's often difficult to ascertain if the cause is due to a mistake or malice, but regardless, the immediate effect is usually the same. Therefore, individual presumptions on the reliability of an operator and subjective assessments if an incident is accidental or intentional are not only a regular feature, but a key aspect of routing security.

Solving the problem of BGP insecurity to prevent future route hijacks and even route leaks requires considerable coordination in the Internet community, a concept that fundamentally goes against the distributed action and autonomous operations design tenets of BGP. Once an ISP or end site receives an internet address block and an AS number allocation from its Regional Internet Registry (RIR),[16] typically, it would need to register specific information to include the local "origin AS" / IP address block (prefix) associations in one or more Internet Routing Registries (IRRs)[17] so that their ISPs and potentially other networks can generate routing policies and "filters" in accordance with local policies. Furthermore, each network may be required by its ISP to publish routing policies regarding what upstream networks (ASes) are authorized to provide "transit services (e.g., an ISP providing an enterprise global connectivity) for the network's destinations (i.e., authorized upstream peers). Requirements for publication of this information in one or more IRRs is voluntary and is solely up to each individual AS, some of which may proxy register routes in IRRs for their customers, utilize alternative internal customer configuration and routing policy databases, or perhaps not require any route registration at all. A key characteristic of the BGP system is that any AS can potentially announce reachability for any IP addresses to the entire world, meaning that any single AS can potentially have a detrimental effect on the global reachability of any Internet destination.

For instance, if the routing information is published in an IRR, other non-adjacent network operators may also use that information to provision routing policies in their routers. The complexity of computing IRR-derived filters for each feasible path to reach a given destination can be considerable for large network operators, especially as new networks and network interconnections are added and as one moves closer to the largest "tier-1" networks at the core of the Internet, where even the largest routers today can't load policy information for all the feasible paths reachable via each of its BGP neighbors. Routing policies may specify whether to accept one or more specific routes from one or more peers and/or customers, and, with a specific origin AS, from a particular peer that has been authorized to announce the route.

RFC 7682[18] outlines some of the historical and existing challenges with the IRR model. The most significant of these challenges is that there are a multitude of IRRs in operation,[19] some operated by ISPs, some operated by research and academic institutions, some by RIRs,[20] and some by for-profit entities. With a few exceptions (e.g., RIPE IRR[21]), there is little to no strict tethering of who holds what number resources with who is authorized to publish routing information for those number resources in any given IRR. As a result, bad actors, misconfigurations, automated proxy registrations by ISPs, or other errors have resulted in a large amount of information being published in IRRs that may not be reliable for provisioning of inter-domain routing policies and may even cause unintended scaling or security issues. Furthermore, the data stored and provided by IRRs is not cryptographically verifiable by relying parties, and stale information is rarely purged from the IRR system. Despite these shortcomings, most inter-domain routing policies today are still provisioned based on the IRR system.

## The Resource Public Key Infrastructure (RPKI)

Fortunately, there is a solution already available and gaining considerable deployment traction. A new system, primarily supported by the five RIRs,[22] is referred to as Resource PKI (RPKI) and provides a cryptographic number resource certification infrastructure. The RPKI enables Internet number allocation authorities and resource holders (e.g., ISPs and end sites) to specify "Route Origin Authorizations (ROAs)" that are cryptographically verifiable and can be used by relying parties (i.e., network operators) for ingesting route origin verification data. That data can be used to automate ingestion of data and configuration of origin validation routing policies directly into routers, automating much of what were historically cumbersome workloads that were prone to operational issues and configuration "drift," and complex for even the most sophisticated routers to process. This nascent RPKI system was developed in the IETF and is standards-based. The RPKI does appear to be gaining traction[23,24] and will certainly address many of the issues that led to decay of various sorts with the current IRR system. Furthermore, it could also be used to bootstrap or otherwise inform and revitalize the IRRs, allowing network operators to identify what information in an IRR was derived from the RPKI and which can therefore be cryptographically validated and associated with routing policies.

The RPKI does appear to be gaining traction, and will certainly address many of the issues that led to decay of various sorts with the current IRR system.

RPKI brings a new set of challenges of its own. Foremost, RPKI creates new external and third-party dependencies that, as adoption continues, ultimately challenge the autonomous operations of the routing system and, if too tightly coupled to the routing system, may impact the robustness and resilience of the Internet itself. RPKI relies on the DNS, and the DNS depends on the routing system. Therefore, particular attention needs to be paid to these interdependencies. Specifically, with RPKI, network operators need to be careful not to introduce tightly coupled circular dependences where the routing system in turn relies on the RPKI, especially at times of startup and instability, otherwise recovering from instability and outages could result in race conditions (i.e., where a system tries to perform multiple functions in parallel that need to be done in sequence[25]) or other bootstrapping issues. This threat can be avoided by ensuring proper operational buffers are in place to absorb failures to various components of the system. A great deal of research has been done considering systemic dependencies and their implications on communications resilience (e.g., the NSTAC Report to the President on Communications Resiliency[26]), and the RPKI system itself would certainly fall into this category of "public core of the Internet" and should be factored into account accordingly.

Perhaps the most significant challenge to RPKI is how the activities of the RIRs can potentially have direct operational implications on the routing system. Unlike the DNS, the global RPKI as deployed does not cleanly model the number resource allocation hierarchy and does not have a single root. Instead, it has multiple trust anchors, operated by each of the RIRs. Currently, the RIRs "over-claim" number resources[27,28,29] to ease complexity of number resource transfers between RIRs.[30] This effectively puts the onus on the relying parties (i.e., network operators)[31] to resolve conflicts should they occur, whereas those relying parties have little to no capability to resolve such conflicts (i.e., how could they know which of two remote ASes that received number resources from different RIRs is the authorized entity to originate a given route?). It also means that a compromise of any RIR's RPKI infrastructure could potentially impact the entire system—regardless of from where a number resource was assigned. While one potential mitigating control is for RIRs to greatly increase the security and stability of their RPKI infrastructure, they'll still be prone to attacks and operational errors alike. If the RIRs were to refrain from overclaiming number resources (and address

the transfer issues via other means), then operators would need to worry primarily about their RIR as far as routing of the prefixes they originate goes. Each Operator would need to interface with all of the RPKI infrastructure when they develop and generate their own routing policies. Even then, a fully operational RPKI that's used to develop routing policy by network operators more broadly will require the RIRs to develop and maintain levels of security and 24x7 operations for which they've traditionally not been funded or required to provide,[32] a growing pain their members are surely going to need to fund in the coming years.

While a cryptographically verifiable number resource allocation repository is a necessity for securing the routing system,[33] just how loosely or tightly coupled that system is to the current Internet routing system will ultimately determine the fragility of the system, and the ability for entities of that system to preserve necessary autonomy in operations. Furthermore, by the very nature of bolting a hierarchical system on to a loosely distributed routing system, the RPKI itself potentially introduces new control points (e.g., the RIRs themselves) and security vulnerabilities. These include so-called "grandparenting" attacks (where someone in the allocation hierarchy takes an action undesirable to the resource holder)[34] and other attacks that may not necessarily exist in the inherently insecure and loosely coupled legacy IRR model, where routing by rumor is the norm. The ideal state is to find a balance between the vital new structure of the RPKI, as well as the inherently ad hoc but tried-and-tested system of routing by rumor.

The collection of systems that makeup the RPKI is very nascent. The scale, stability, and security of the RIR infrastructure that constitute much of the RPKI will play a much more critical role in the operations and security of the routing system in the future than RIR systems have historically played. Traditionally, RIRs allocate Internet number resources (address space and AS numbers) to ISPs or end sites and make available information associated with those allocations via WHOIS[35] or other means. Beyond perhaps operating various components of DNS infrastructure and an IRR themselves, RIRs had no direct operational tie-in to how the number resources are utilized in the routing system. An RPKI-enabled routing system requires constant maintenance, high performance, robust security, and high availability. This is a significant departure from the traditional operational expectations of RIRs. The increased operational importance of RIRs means that they, too, should be considered part of the public core of the Internet.

Given the risks associated with this new role, the RIRs are still evolving their own organizational thinking, from both legal[36] and technical perspectives,[37] and are prudently reminding relying parties to be cautious when coupling the RPKI to their network routing policies[38] without sufficient operational buffers.

With the growing reliance on the Internet for mission-critical functions, and continuing concern about insecurities of the routing system, the promise of RPKI to ameliorate some of the vulnerabilities is being well received, as evidenced by its rate of adoption. RPKI has seen significant growth in adoption over the last three years, from ~10% of registered Prefix-Origin pairs having RPKI validation data at the end of 2018, to ~31% valid Prefix-Origin pairs in October 2021. RPKI adoption percentages are not uniform at each RIR, yet by any measure, they've been impressive.[39]

Beyond the RIRs, a significant number of Tier-1 telecom providers (e.g., GTT, NTT, and Telia[40]) and large network operators have fully implemented RPKI-based origin validation. According to "Is BGP Safe Yet"[41] (which conjectures that RPKI makes it safe), 102 known operators worldwide have completed the full implementation of RPKI. An additional 24 operators have partial RPKI deployment, and another 240 operators have only just begun the process of RPKI deployment. While that

still leaves another ~72,000 networks[42] to act, it is a significant deployment rate in such a relatively short timeframe, especially when compared to historic IPv6 and DNSSEC deployment rates.

The original objective of RPKI-based origin validation was to prevent perhaps the most significant class of notable routing security incidents, those that involve re-origination of routes the local AS is not authorized to announce. Re-origination incidents are commonly the result of router policy misconfiguration or buggy software. The Pakistan / YouTube Incident,[43] in which Pakistan announced YouTube address space globally while attempting to censor it locally within Pakistan, had the effect of taking all of YouTube offline globally. Another similar incident, , commonly referred to as the infamous AS7007 incident,[44] occurred when a BGP router operated by AS7007 accidently announced to the Internet that it was the proper destination AS for a large portion of Internet address space. In these types of incidents, mis-origination was easily identifiable. The impact with these and similar incidents is commonly compounded when the routing announcements are "more specific" than the legitimate announcements, as IP routing protocols normally always prefer the most specific route over less specific routes.

Without RPKI, a sophisticated attacker can likely circumvent AS origin validation alone quite easily, and it commonly happens even by default with many forms of route leaks, although when it does, it makes intent of the misbehaving network easier to identify. Origin validation, be it based on RPKI or IRR routing policy information, will certainly prevent an entire class of BGP security incidents that occur commonly today.

One final note with RPKI-based origin validation and its IRR-based counterparts, however, is that the manipulation of a BGP AS path, including the origin, is still possible, and until cryptographic security protocols that link RPKI to routing protocol integrity protections can be deployed at scale, this problem will persist. There has been a large amount of additional work to address BGP path validation beyond just the origin via protocols such as BGPsec,[45] where RPKI-derived cryptographic signatures are attached to information within the routing system and BGP itself to provide integrity protections. However, it remains to be seen if this work is worth the complexity and fragility it introduces, especially as it is still vulnerable to route leaks[46] and other similar forms of attack that need to be addressed via peering and operational best practices, as discussed in the next section.

While attacks leveraging the routing system can be targeted and intentionally scoped, most attacks in the routing system are noisy, globally propagated, and fairly trivial to detect. Of course, as discussed previously, discerning whether a given routing security incident was the result of malice or error is complex for external observers. While the immediate effect is often the same, it is common that little to no authoritative information on the root cause for a given incident ever emerges. Fortunately, network operators can take decisive action to filter and "reverse" bad routing information once identified, and the offending network(s) are commonly identified in operational and security forums, but there is often little to no recourse. This noise factor associated with routing system attacks is likely an attribute from where much restraint for launching such attacks stems, for state and non-state actors alike. Yet they still seem to have occurred,[47,48,49,50] and likely will continue to occur, and even if only temporary, it's important to recognize that the attacker's objective may have been achieved. As with most security, this is where layered defenses and best practices come into play, as discussed in the following section.

Beyond the supporting infrastructure mechanisms (e.g., RPKI and IRRs) noted above, there is an Internet Society[51] initiative that focuses on Mutually Agreed Norms for Routing Security (MANRS),[52] which aims to help reduce the most common routing system vulnerabilities. The objective of

MANRS is to improve the security and reliability of the global Internet routing system, based on collaboration among participants and shared responsibility for Internet infrastructure, setting a new norm for routing security and network operators. The specific categories of participants in MANRS include network operators and Internet exchange points, as well as content delivery networks and cloud providers. The MANRS program aims to raise awareness and create a culture of collective responsibility toward the security and resilience of the global routing system. It does this by providing a framework for network operators to better understand and address issues relating to the security and resilience of the routing system, to include best practices to prevent propagation of incorrect routing information, preventing traffic with spoofed source IP addresses, facilitating communication among operations, facilitating publication and validation of routing information on a global scale, and providing monitoring and debugging tools to participants. The MANRS initiative is continuing to gain traction and is certainly helping to make the routing system more secure.

## Conclusion

While there are a broad array of other considerations related to attacks against the routing system, increased tooling and infrastructure to address the threat posed by route leaks and route hijacks will surely go a long way toward better securing the routing system. A stable and secure cryptographic number resource certification infrastructure is an absolute necessity to inform routing policies used to secure the routing system. However, the Internet community must be cautious to understand the implications of introducing potential new control points and systemic dependencies—and how they may impact the resilience, flexibility, and autonomy in operations for each participating network—that have made the current routing system so robust and successful.

Routing by rumor has served us well, and a decade ago it may have been ideal because it avoids systemic dependencies—but it is certainly past its prime in today's cyber environment. The accumulated improvements discussed here and elsewhere are changing rumors into knowledge and will ideally provide the foundation for a more secure Internet routing system in the future.

Currently, some of the discussion around the application of the public core definition within routing has focused on the importance of addressing routing hijacks, such as those discussed above. These remain difficult to address if intentional and launched by a sophisticated adversary in cooperation with one or more network operators. However, one category of routing incidents has been a key focus thus far, and for this RPKI will help significantly. However, this solution also increases the operational importance of previously less relevant organizations (i.e., the RIRs) and the infrastructure they operate. This change and its ramifications must be fully understood and considered by all stakeholders (to include the memberships of the RIRs), given the full set of new obligations and resource allocation requirements that has been placed upon them. Together, however, such improvements represent a welcome maturation of the routing system away from just "routing by rumor" to "routing by fact."

Currently, some of the discussion around the application of the public core definition within routing has focused on the importance of addressing routing hijacks. These remain difficult to address if intentional and launched by a sophisticated adversary in cooperation with one or more network operators.

## Endnotes


1 Global Commission on the Stability of Cyberspace, Advancing Cyberstability, The Hague: The GCSC, November 2019. https://cyberstability.org/wp-content/uploads/2020/02/GCSC-Advancing-Cyberstability.pdf

2 "Norm on the Non-Interference with the Public Core," Global Commission on the Stability of Cyberspace, https://cyberstability.org/norms/#toggle-id-1

3 "Cyberstability Paper Series. New Conditions and Constellations in Cyber," Global Commission on the Stability of Cyberspace, https://cyberstability.org/paper-series/

4 Each of the five RIRs serve different regions. AFRINIC serves Africa and portions of the Indian Ocean, APNIC serves portions of Asia and portions of Oceania, ARIN serves Canada, many Caribbean and North Atlantic Islands, and the United States, LACNIC serves Latin America and portions of the Caribbean, and RIPE NCC serves Europe, the Middle East, and Central Asia.

5 "The Internet Engineering Task Force", https://ietf.org

6 "The CIDR Report", https://www.cidr-report.org/as2.0/

7 "The Interconnection Database", PeeringDB, https://www.peeringdb.com

8 "The CIDR Report", https://www.cidr-report.org

9 IEFT, "A Border Gateway Protocol 4 (BGP-4)," IETF Datatracker, RFC 4271, https://datatracker.ietf.org/doc/html/rfc4271

10 William Jr Haag, Doug Montgomery, William C. Baker and Allen Tan, Secure Inter-Domain Routing, (National Institute of Standards and Technology and National Cybersecurity Center of Excellence, 2016). https://www.nccoe.nist.gov/sites/default/files/library/project-descriptions/sidr-project-description-draft.pdf

11 K. Sriram, et al. "Problem Definition and Classification of BGP Route Leaks," Internet Engineering Task Force, RFC7908, (June 2016) https://www.rfc-editor.org/rfc/rfc7908.txt

12 Masahito Ando, Masayuki Okada, and Akira Kanaoka, "Simulation Study of BGP Origin Validation Effect against Mis-Origination with Internet Topology," 12th Asia Joint Conference on Information Security, (August 2017): doi 10.1109/AsiaJCIS.2017.17. https://www.researchgate.net/publication/319587923_Simulation_Study_of_BGP_Origin_Validation_Effect_against_Mis-Origination_with_Internet_Topology

13 Ryan Singel, "Pakistan's Accidental YouTube Re-Routing Exposes Trust Flaw in Net," Wired, February 25, 2008, https://www.wired.com/2008/02/pakistans-accid/

14 Ax Sharma, "Major BGP Leak Disrupts Thousands of Networks Globally," Bleeping Computer, April 17, 2021 https://www.bleepingcomputer.com/news/security/major-bgp-leak-disrupts-thousands-of-networks-globally/

15 "BGP Routing Leak Detection System", https://puck.nether.net/bgp/leakinfo.cgi

16 "Regional Internet Registries", The Number Resource Organization, https://www.nro.net/about/rirs/

17 "Internet Reouting Registry", http://www.irr.net

18 Danny McPherson, et al. "Considerations for Internet Routing Registries and Routing Policy Configuration", Internet Engineering Task Force, RFC 7682 (December 2015), https://tools.ietf.org/html/rfc7682

19 "List of Routing Registries", Internet Routing Registry, http://www.irr.net/docs/list.html

20 EXPLAIN RIRs

21 "Managing Route Objects in the IRR", RIPE Network Coordination Centre, https://www.ripe.net/manage-ips-and-asns/db/support/managing-route-objects-in-the-irr

22 "Contact Your Local Regional Internet Registry," The Number Resource Organization, https://www.nro.net

23 "NIST RPKI Monitor", National Institute of Standards and Technology, https://rpki-moni-

tor.antd.nist.gov

24 "Is BGP Safe Yet? No.", https://isbgpsafeyet.com

25 Ben Lutkevich, and Brien Posey, "What is a race condition?", TechTarget, last modified June 2021, https://searchstorage.techtarget.com/definition/race-condition

26 The President's National Security Telecommunications Advisory Committee, "NSTAC Report to the President on Communications Resiliency," CISA, May 2021, https://www.cisa.gov/sites/default/files/publications/NSTAC%20Report%20to%20the%20President%20on%20Communications%20Resiliency.pdf

27 "Regional Internet Registries are preparing to deploy "All Resources" RPKI Service," The Number Resource Organization, July 11, 2017, https://www.nro.net/regional-internet-registries-are-preparing-to-deploy-all-resources-rpki-service/

28 A. Newton et al., "RPKI Multiple "All Resources" Trust Anchor Applicability Statement," Internet Engineering Task Force, (July 2016), https://www.ietf.org/archive/id/draft-rir-rpki-allresta-app-statement-01.txt

29 Geoff Huston, "RPKI and Trust Anchors," APNIC, April 21, 2020, https://blog.apnic.net/2020/04/21/rpki-and-trust-anchors/

30 "IAB Statement on RPKI," Internet Architecture Board, April 3, 2018, https://www.iab.org/documents/correspondence-reports-documents/2018-2/iab-statement-on-the-rpki/

31 "IAB Statement on RPKI," Internet Architecture Board, January 27, 2010, https://www.iab.org/documents/correspondence-reports-documents/docs2010/iab-statement-on-the-rpki/

32 An illegitimate RPKI ROA could result in origin validation actions that suppress legitimate routing system announcements within ISP networks that are performing origin validation filtering. This issue exists whether or not a network's number resources were allocated from that specific RIR and even whether or not the network is currently using RPKI to publish ROAs. What specific route filtering actions a network operator might take based on the presence of a valid ROA that matches a route received from a BGP peer, or an invalid ROA, or no matching or "covering" ROA at all, is solely dependent on how RPKI-derived enforcement policies are currently implemented by each of the individual network operators in the routing system. These policy enforcement actions independently effectuated by each ISP may range from suppressing or completely filtering a route, to lowering or increasing the preference, or they may take no action at all.

33 Martin J. Levi, "RPKI—The required cryptographic upgrade to BGP routing," Cloudfare, September 19, 2018, https://blog.cloudflare.com/rpki/

34 Danny Cooper, et al. "On the Risk of Misbehaving RPKI Authorities," Boston University, Department of Computer Science, (November 2013) https://www.cs.bu.edu/fac/goldbe/papers/hotRPKI.pdf

35 "What is a Whois Service?," American Registry for Internet Numbers, https://www.arin.net/resources/registry/whois/

36 "ARIN's Trust Anchor Locator (TAL). Relying Party Agreement (RPA)," American Registry for Internet Numbers, https://www.arin.net/resources/manage/rpki/tal/

37 Felipe Victolla Silveira, "RIPE NCC and the Cloud: Draft Principles, Requirements and Strategy Framework," RIPE Labs, August 3, 2021, https://labs.ripe.net/author/felipe_victolla_silveira/ripe-ncc-and-the-cloud-draft-principles-requirements-and-strategy-framework/

38 "Notice of upcoming maintenance to ARIN's RPKI Infrastructure," American Registry for Internet Numbers, June 2, 2021, https://www.arin.net/announcements/20210602-rpki/

39 Where uptake has been slowest (for instance, across Africa), external support would be crucial, perhaps within currently envisioned state-led capacity-building (development assistance) programs being developed. Indeed, the crucial operations and needs of RIRs seem to have been somewhat ignored in the current debate on capacity building, and this example illustrates where additional resources could be meaningfully directed by the global policy community: RIPE, 47%

Valid; LACNIC, 37% Valid; APNIC, 34% Valid; ARIN, 20% Valid; AFRNIC, 13% Valid.

40 Louis Poinsignon, "The Internet is Getting Safer: Fall 2020 RPKI Update," Cloudfare, June 6, 2021, https://blog.cloudflare.com/rpki-2020-fall-update/

41 "Is BGP Safe Yet? No.", https://isbgpsafeyet.com

42 "The CIDR Report", https://www.cidr-report.org

43 Ryan Singel, "Pakistan's Accidental YouTube Re-Routing Exposes Trust Flaw in Net," Wired, February 25, 2008 https://www.wired.com/2008/02/pakistans-accid/

44 "BGP Case Studies," BGP.us, https://www.bgp.us/case-studies/

45 M. Lepinski, and K. Sriram, "BGPSEC Protocol Specification," Internet Engineering Task Force, RFC 8205, (September 2017) https://datatracker.ietf.org/doc/html/rfc8205

46 Danny McPherson et al. "Route -Leaks & MITM Attacks Against BGPSEC," Internet-Draft. IETF, (April 2014) https://datatracker.ietf.org/doc/html/draft-ietf-grow-simple-leak-attack-bgpsec-no-help

47 Catalin Cimpanu, "For two hours, a large chunk of European mobile traffic was routed through China," ZDNet, June 7, 2-19, https://www.zdnet.com/article/for-two-hours-a-large-chunk-of-european-mobile-traffic-was-rerouted-through-china/

48 Zak Doffman, "Russia and China 'Hijack' Your Internet Traffic: Here's What You Do," Forbes, April 18, 2020, https://www.forbes.com/sites/zakdoffman/2020/04/18/russia-and-china-behind-internet-hijack-risk-heres-how-to-check-youre-now-secure/?sh=7849a2a65b16

49 Andy Greenberg, "China Hijack 15% of Internet Traffic? More Like .015%," Forbes, November 19, 2010, https://www.forbes.com/sites/andygreenberg/2010/11/19/china-hijacks-15-of-internet-traffic-more-like-015/?sh=56eb0299223a

50 Dan Goodin, "Russian-controlled telecom hijacks financial services' Internet traffic," ArsTechnica, April 27, 2017, https://arstechnica.com/information-technology/2017/04/russian-controlled-telecom-hijacks-financial-services-internet-traffic/

51 "Internet Society," https://www.internetsociety.org

52 "MANRS," https://manrs.org/

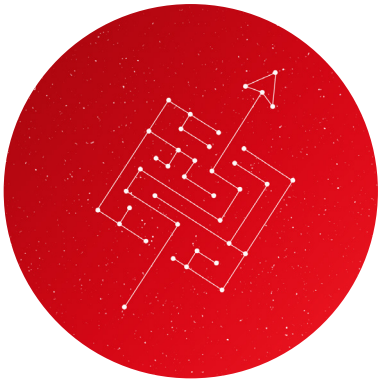

## About the Author

Danny McPherson is responsible for Verisign's information systems and services, as well as information and corporate security. He has actively participated in internet operations, research, and standardization since the early 1990s, to include serving on the Internet Architecture Board (IAB) and chairing an array of Internet Engineering Task Force (IETF) and other working groups and committees. He has authored several books, numerous internet protocol standards, network and security research papers, and other publications.

Previously, McPherson was CSO at Arbor Networks, where he developed solutions to detect and mitigate cyberattacks and performed pioneering research on Internet infrastructure evolution, as well as botnet and malware collection and analysis. Before that, he held technical leadership positions in architecture, engineering and operations with Amber Networks, Qwest Communications, Genuity, MCI Communications, and the U.S. Army Signal Corps.

## About the Cyberstability Paper Series

Since the release of the final report of the Global Commission on the Stability of Cyberspace in November 2019, the concept of cyberstability has continued to evolve. A number of new 'conditions' are emerging: new agreements on norms, capacity building and other stability measures have been proposed and solidified within the United Nations and elsewhere, and stakeholders are exploring ways to increase stability and minimize the risk of conflict in cyberspace through technical fixes or governance structures. The constellations of initiatives involved in working towards cyberstability is expanding, underlining the need to connect the traditional state-led dialogues with those of the Internet communities from civil society and industry. Gaps continue to close, between the global north and south, between technology and policy, but also the stability in and the stability of cyberspace.

The first Cyberstability Paper Series explores these "New Conditions and Constellations in Cyber" by collecting twelve papers from leading experts, each providing a glance into past or future challenges and contributions to cyberstability. The papers are released on a rolling basis from July until December 2021, culminating in an edited volume. All papers will be available for open access, and a limited number of printed hardback copies are available.

## Published by

GLOBAL COMMISSION
ON THE STABILITY OF CYBERSPACE

The Hague Centre
for Strategic Studies

The opinions expressed in this publication are those solely of the author(s) and do not reflect the views of the Global Commission on the Stability of Cyberspace (GCSC), its partners, or The Hague Centre for Strategic Studies (HCSS).

© 2021 The Hague Centre for Strategic Studies and the Global Commission on the Stability of Cyberspace. This work is licensed under a Creative Commons Attribution – Noncommercial – No Derivatives License. To view this license, visit www.creativecommons.org/licenses/by-nc-nd/3.0. For re-use or distribution, please include this copyright notice.